Perception, 2000, volume 29, pages 171 – 200

DOI:10.1068/p3013

Perceptual assumptions and projective distortions in a three-dimensional shape illusion

A Fuzz Griffiths, Qasim Zaidi
State University of New York, State College of Optometry, 100 East 24th Street, New York, NY 10010, USA; e-mail: fuzz@sunyopt.edu
Received 19 November 1998, in revised form 29 July 1999

Abstract. We examine a shape illusion, in which the balconies of a building appear to tilt up or down, depending on the viewpoint. The balconies are actually level parallelogram shapes, but appear as tilted rectangles. We measured the illusory tilts observed when parallelogram shapes are viewed above the line of sight, using three-dimensional stimuli consisting of parallelograms of various tilts viewed at different orientations. Under perspective projection, parallelism and orthogonality are not preserved. However, perspective distortions alone cannot account for the perceived tilts measured in these experiments, since observers perceived illusory tilts even for stimuli in the frontoparallel plane. We introduce a model, based on the theory that observers assume ambiguously projected three-dimensional angles to be equal to 90°, but revise their predictions on the basis of observation. In the model, perceived tilt is predicted as a weighted sum of the tilts predicted by the assumptions that the shape is rectangular, and that the shape is level (ie that the angle between the shape and the vertical backboard is equal to 90°). We prove that it is mathematically impossible for a planar rectangle to share a projection with a nonrectangular parallelogram. A less restrictive assumption that just the two leading internal angles are equal to 90° is suggested as an alternative, and it is further proven that this new configuration of angles leads to a unique perceived tilt. The relative weights in the model reflect the amount that each prediction is revised, and are shown to vary systematically with stimulus orientation. For some observers a better fit was found by replacing the level-tilt assumptions with an assumption that physical tilt was equal to the projected tilt.

1 Introduction

As one views the building shown in figure 1a (Halper 1997), a curious thing happens. From street level, the balconies on the side of the building appear to be tilted up; the further up the building you look, the stronger the appearance of tilt. When viewed from the opposite side (figure 1b), the balconies appear to be tilted downwards. When viewed from a vantage point directly opposite the building, however, the balconies appear to be horizontal (which they are). The illusion is persistent, despite repeated viewing from different angles, as well as the cognitive implausibility of nonrigid, non-level balconies. The balconies in the figure are actually nonrectangular parallelograms. This illusion does not seem to occur in buildings with rectangular balconies, even when the balconies are seen from a variety of different viewpoints.

This is a novel example of a class of illusions in which a false three-dimensional (3-D) percept is inferred from retinal projections (Ames 1951; Ittelson 1952). Related illusions have been reported for other nonrectangular buildings such as the Hancock tower in Boston (Perkins 1976). These illusions can help answer the fundamental question of visual perception. Namely, they can tell us why some things appear the way they do, and why some things appear the way they are (Koffka 1963). In this study, we wish to highlight the mechanisms used to infer shape. By isolating conditions under which shape inferences are inaccurate, we intend to identify some of the assumptions made by the visual system in response to these kinds of stimuli.

Under perspective projection, properties such as parallelism and orthogonality remain intact only if they lie frontoparallel to the image plane on the light of sight. Otherwise, distortions occur, and angles in the image are no longer equal to the 3-D angles in

(a) (b)

Figure 1. (a) A view of the apartment building 'The Future', located at 200 East 32nd Street in New York City. The balconies appear to tilt up, with tilt becoming more apparent the higher up the building one looks. (b) The same building as that shown at left, but viewed from the opposite side. The balconies now appear to tilt down.

the real world. The images of points lying on planes that are not frontoparallel to the image plane are shifted toward the vanishing point as a function of their distance from the image plane. As a result, a solid 3-D angle of 90° can project to an angle in the image plane which is greater than, equal to, or less than 90°, depending on the orientation of the lines relative to the image plane and line of sight. For example, in the retinal projection of a straight edge viewed at an oblique angle to the image plane, the end farther from the eye is projected closer to the vanishing point (ie tilted down if viewed from below).

The transformation of an image under perspective projection does not preserve local angles between lines or edges, and so we refer to this loss as 'projective distortion'. The image of the balconies in figure 1, which appear to be angled up or down, is one example of a projective distortion. However, rectangular balconies suffer the same projective distortions when viewed from an oblique angle, yet do not generally appear physically tilted to the observer.

To test possible explanations for the illusion described above, we used a large set of physical 3-D models, such as those shown in figure 2. In figure 2a, the image is a photograph of a horizontal parallelogram object with internal 2-D shape angles equal to 135°/45°. This parallelogram is mounted against a vertical backboard and presented so that the backboard is along the line of sight. The short edge of the parallelogram, closest to the observer (hereafter referred to as the 'leading edge') is oriented at an angle, in the $x-z$ plane (equivalent to an overhead view, where x represents lateral displacement, y represents height, and z represents depth, relative to the observer), of

45° to the line of sight. For the purposes of this paper, the 'internal angle', θ, of a shape is defined as the 2-D angle between the leading edge of the parallelogram and the parallelogram's long edge along the backboard (see figure 3a). The variable 'orientation', ω, of the leading (shorter) edge is defined relative to the line of sight in the $x-z$ plane. 'Tilt', ψ, will be used to describe the physical orientation of a surface relative to the horizontal, independent of viewpoint (see figure 3b). All the symbols used in this paper are defined in table 1.

Table 1. Definitions of symbols used in the paper.

Symbol	Definition
ψ	physical tilt of the stimulus with respect to horizontal axis
θ	internal angle of the parallelogram
τ	projected 2-D tilt of the leading (shorter) edge
ρ	perceived 3-D tilt of the leading (shorter) edge
v	3-D tilt of the leading edge of the virtual rectangle
ω	angle between line of sight and the leading edge (viewing angle)

The image in figure 2b is a photograph of a horizontal ($\psi = 0°$) rectangular object on a vertical backboard, presented so that the leading edge has the same orientation ($\omega = 45°$) as the leading edge of the object in figure 2a. Most observers see the parallelogram object ($\theta = 135°$) as tilted up, and the rectangular object as level. Measurement of the 2-D angles in the photographs reveals that the projected angle between the vertical backboard and the leading edge is the same (65°) in both photographs.

We chose to use physical 3-D stimuli rather than their 2-D images for two principal reasons. First, we are concerned with the projection of a physical object, and the ways in which observers make use of the information contained therein. Our stimuli were presented above the eye level of observers. If one were to present images taken

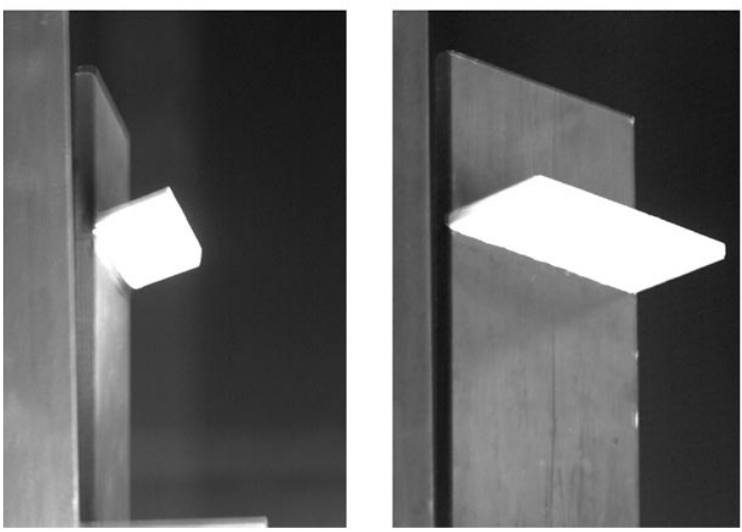

(a) (b)

Figure 2. (a) A solid parallelogram shape with internal angle, θ, equal to 135°, viewed from below with the leading edge oriented 45° towards the camera from line of sight ($\omega = 45°$). There is no physical tilt in this stimulus, but observers typically perceive this shape as tilted up, relative to the horizontal. (b) A rectangle ($\theta = 90°$) viewed with similar orientation, that is with the leading edge oriented 45° toward the camera from line of sight ($\omega = 45°$). Once again, there is no physical tilt. Observers typically perceive this stimulus veridically as horizontal.

from a similar viewpoint, effects would depend on where one chooses to present these images relative to the observer. Presenting these images at eye level in the fronto-parallel plane gives a false impression of the orientation of stimuli relative to the observer, whereas presenting images in other locations introduces a new set of projective distortions of the images. Under perspective projection, the projection of a projected image is not the same as the projected image itself (Mundy and Zisserman 1992).

Second, the use of 2-D images, rather than 3-D stimuli, presents observers with two possible ways in which to evaluate aspects of the stimuli, eg tilt of an edge. Tilts can be estimated relative to other parts of the scene (other edges), as represented in the image. Alternatively, tilts can be evaluated as the orientation of a projected edge relative to the surface of the image plane itself, eg the angle between the projected edge and the borders of a photograph. This 'crosstalk' between the surface and scene properties of an image has been shown to introduce biases into empirical studies of this nature (Sedgwick and Nicholls 1993; Sedgwick et al 1995). In figure 2, both of the stimulus shapes are mounted at 90° to the backboard (they are physically level), but both projections in the image of the leading edges are angled up relative to the horizontal.

In this paper, we investigate this illusion in order to identify those properties of the scene and the visual system that combine to yield illusory percepts. In the first experiment, we replicated the original building illusion in the laboratory under controlled conditions, and empirically determined the size of the effect. We speculated that the visual system may be unable to account for the projective transformation by which the image is formed. In the second experiment, we presented the stimulus shapes with their leading (shorter) edges at 90° to the observer's line of sight in the $x-z$ plane. In doing this, we equated the projections in the image plane of all the leading edges of a given tilt. This manipulation allowed us to control for factors other than shape, since the projected angle in the image plane was equivalent to the physical tilt of the parallelogram shape for level stimuli ($\psi = 0°$). (For stimuli with nonzero tilts, the projected angle differed from the physical tilt by less than 2°.)

In the third experiment, we presented the stimuli with their leading edges at a constant angle (45° or 135°) to the line of sight in the $x-z$ plane. Under these conditions, the projections of the leading edges were constant across stimuli of different shapes, but the projected angle in the image plane was different from the physical tilt by approximately 30°. In doing so, we manipulated the stimuli so that the angles in the physical world were no longer equal to those in the projected image, but this difference was held constant across all stimulus shapes (as opposed to the first experiment, where the difference was a function of internal angle). In a fourth experiment we investigated whether observers were able to accurately estimate the shapes of the parallelogram stimuli themselves. On the basis of the results of these experiments we developed and tested linear models of cue combination in order to determine whether prior assumptions or perceptual tendencies within the visual system play a role in the illusion.

2 Methods

2.1 *Subjects*

We report data from four observers, all of whom had corrected-to-normal vision. The observers had previously seen the building shown in figure 1, and had experienced the associated illusion first-hand, but were unaware of possible explanations for the illusion. Three sets of measurements were made for all conditions on two observers (AL and SB). Single sets of measurements were also made on two other observers (JM and MA) for experiments 2 through 4. The data for JM and MA are shown only in the modeling sections.

2.2 *Stimuli*

The test stimuli used to simulate the balconies in the building consisted of a set of seven 4 inches × 2 inches parallelogram-shaped objects (see figure 3a). The internal angle, θ, ranged from 45° to 135°, in 15° increments. These shapes were constructed out of $\frac{3}{8}$ inch thick balsa wood, painted white and mounted along the longer side against a vertical black backboard 4 inches wide, as shown in figure 2. Stimuli were mounted either level, ie perpendicular to the backboard, or were tilted up/down by 15°. This physical tilt relative to the horizontal is shown as ψ in figure 3b. Positive values of ψ are used to represent upward tilt, while negative values represent downward tilt. The seven shapes ($\theta = 45°$, 60°, 75°, 90°, 105°, 120°, 135°) and three physical tilts ($\psi = -15°$, 0°, +15°) were combined to yield a set of 21 stimuli, all of which were used in each of the experiments.

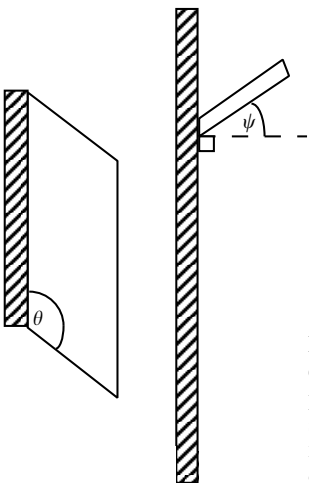

Figure 3. (a) Top view of a typical stimulus shape used in these experiments. The internal angle is represented by θ. The shaded portion represents the backboard. (b) Side view of the stimulus, showing the physical tilt, ψ. Positive values of ψ represent tilts up from the horizontal, as shown here, while negative values represent a downward tilt.

In all the experiments, observers viewed the experimental stimuli monocularly, from a distance of 6 feet. Stimuli were presented at an elevation, ϕ, of 30° above eye level, and were displayed against a black backdrop. A narrow window was constructed by placing two flat black walls on either side of the stimulus, in order to provide a view limited to the visible range of the stimulus shape. Stimuli were illuminated with diffuse overhead light, supplemented by light from a lamp positioned below the stimulus, in order to remove shadows on the stimulus. The supplemental lamp was hidden from the observer. The stimulus was of high contrast, with a matte finish, and shadows and internal reflections were minimized.

2.3 *Procedure*

Observers were asked to estimate the 3-D angle between the parallelogram shape and the backboard, and to select the best representative from a set of 2-D comparison stimuli (see figure 4), presented at eye level. Tilt angles in the comparison stimuli varied between 45° and 135°, in 5° increments. Before commencing the experiment, observers were allowed to view a subset of the experimental stimuli, including representatives of each of the different internal angles, θ, and physical tilts, ψ, used. Observers were told that these were "some of the stimuli we will be using in the experiment". Observers were then trained to estimate angles using the three rectangular stimuli ($\theta = 90°$; $\psi = -15°$, 0°, +15°). During the training phase, observers viewed the stimuli binocularly, close-up, and at eye level, and were asked to estimate the tilt and select the closest comparison stimulus. The observer's response consisted of a code number, corresponding to the

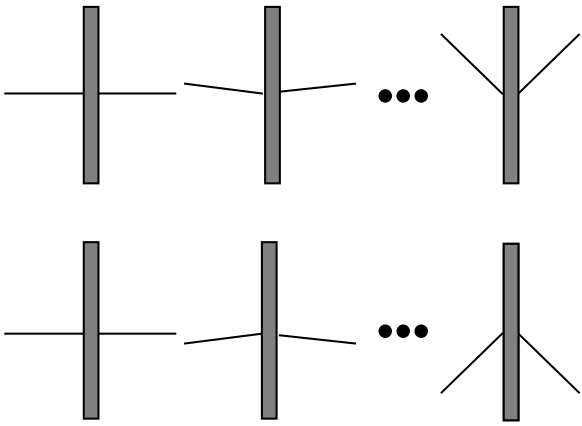

Figure 4. Example set of comparison stimuli used in the tilt-matching experiments. Top row, from left to right, shows tilts, ψ, of 0°, +5°, and +45°; the bottom row, from left to right, shows tilts, ψ, of 0°, −5°, −45°. Tilted edges are shown on either side of each vertical backboard (dark gray shading), since stimuli could be presented with either the left or right side mounted against the backboard. Comparison stimuli were identified by code numbers, which have been omitted from this figure.

comparison stimulus identifier. During training, observers were given accurate feedback. The experimenter took special care to emphasize that observers were being asked to estimate the 3-D tilt angle in the stimuli. This practice session was performed once only. The testing phase commenced after both the observer and the experimenter felt confident that the observer understood the task.

On each trial, observers were able to view the stimulus for a maximum of 10 s, after which they recorded the comparison stimulus identifier on a response sheet. In between trials, the apparatus was hidden from observers while the next stimulus was set into place. No feedback was given during the experiment. Trials from experiments 1 through 3 were interleaved, and presented in random order. A block consisted of approximately 80 trials, and lasted less than 1 h. Observers AL and SB performed three such blocks of trials. Observers JM and MA performed one such block of trials for experiments 2 and 3 only.

The critical variant in each of the following experiments was the orientation of the leading edge of the stimulus shape relative to the observer's line of sight. In experiment 1, the orientation of the leading edge was dependent on the shape of the stimulus, and the orientation of the backboard was parallel to the line of sight. In the other experiments, the orientation of the leading edge was held at a constant angle to the line of sight, while the orientation of the backboard varied as a function of internal angle.

3 Experiment 1

In the first experiment, stimuli were presented such that the backboard was parallel to the observer's line of sight in the $x-z$ plane, and the orientation of the leading (shorter) edge was dependent on the internal angle, so that $\omega = 180^\circ - \theta$ (see figure 5). In this condition, the line of sight was along the backboard and thus parallel (in the $x-z$ plane) to the longer edges of the parallelogram. For tilted parallelograms, the backboard was slightly angled with respect to the line of sight in order to keep ω equal to $180^\circ - \theta$. This deviation was always less than 3°.

Results from this experiment are shown in figure 6. Perceived tilt is plotted as a function of the internal angle, θ, of the parallelogram stimulus. Results from three physical tilts are presented: $\psi = 0^\circ$, $+15^\circ$, -15°. Each data point represents the results from a single trial. Owing to the high repeatability of measurements and the relatively

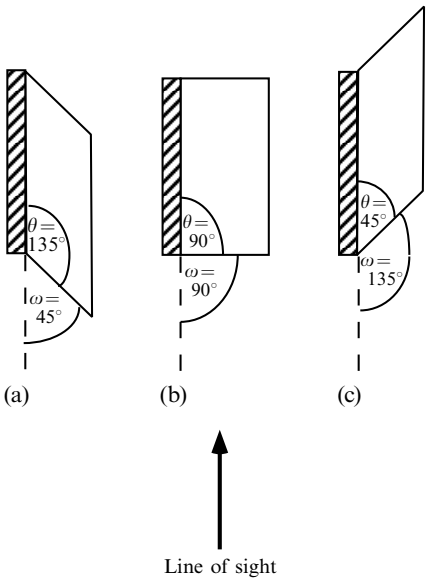

Figure 5. Top view of some of the stimuli presented in the $\omega = 180^\circ - \theta$ viewing condition, with internal angle, θ, equal to: (a) 135°, (b) 90°, and (c) 45°. Line of sight is parallel to the backboard, with observers facing 'due North', towards the top of the page.

coarse quantization of the responses, two or more data points often shared the same value, so that many of the points on the graphs would be superimposed. As a result, the 63 data points would be represented by symbols in only 31 (AL) or 30 (SB) locations. To avoid the occlusion of data points, data from repeated trials have been shifted along the x-axis by $+3^\circ$ or -3°. The three physical tilts, ψ, used are shown as the horizontal dotted lines. For the three different physical tilts, the projected angles between the leading edges and the backboard are shown as dashed curves. Since these angles represent the 2-D projections of a 3-D physical tilt, we have elected to use the term 'projected tilt', to refer to the deviation from the horizontal of the projection of a leading edge in the image plane, which lies orthogonal to the line of sight.

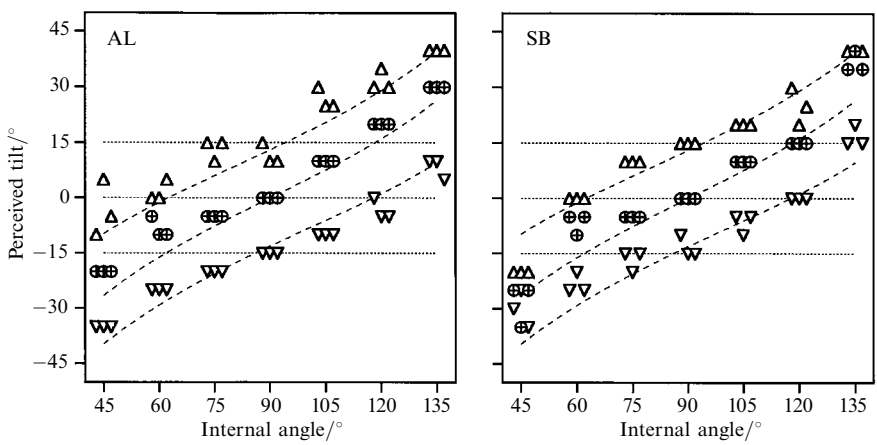

Figure 6. Tilt judgments obtained from experiment 1 ($\omega = 180^\circ - \theta$). Results from three physical tilts are presented, including $\psi = 0^\circ$ (no physical tilt), as shown by the crosses-in-circles, $\psi = +15^\circ$ (upwards tilt), as shown by the upward-pointing triangles, and $\psi = -15^\circ$ (downwards tilt), as shown by the downward-pointing triangles. Three measurements are shown for each value of internal angle for each of the three physical tilts. In order to avoid occlusion, two of the three measurements for each internal angle have been shifted horizontally. The three dashed curves represent the projected tilts of the three physical tilts used, while the physical tilts themselves are represented by dotted lines (see text for details).

Projected tilt, τ, was calculated as a function of the orientation of the leading edge, ω (which is equivalent to $180^\circ - \theta$ under these viewing conditions), physical tilt, ψ, and the angle of elevation of the line of sight, ϕ, by equation (1), derived in Appendix 1.

$$\tau = \arctan\left(\frac{\tan\psi\cos\phi + \cos\omega\sin\phi}{\sin\omega}\right) . \quad (1)$$

Since the leading edges of the level ($\psi = 0^\circ$) stimuli with 90° internal angles (rectangles) lie in the frontoparallel plane, the projected tilts of these edges in the image plane were equal to the physical tilts of the stimuli. For nonlevel ($\psi = +15^\circ$, -15°) rectangular stimuli, projected tilts of the edges are within 2° of the physical tilts. For other (nonrectangular) shapes, the angle between the projections of the backboard and the leading edge in the image plane was greater or smaller than the 3-D tilt angle.

For both observers, perceived tilt was a monotonically increasing function of internal angle (and also as a function of orientation of the leading edge, since $\omega = 180^\circ - \theta$), and there was a clear separation between perceived tilts for the three sets of physical tilts. Perceived tilt was veridical for $\theta = 90^\circ$ shapes, but for other shapes it deviated from the physical tilt, and more closely resembled the projected tilt. Perceived tilts deviated from the projected tilts for the highest and lowest values of internal angle ($\theta = 45^\circ$, 135°). This was especially true for the stimuli with nonzero physical tilt ($\psi = +15^\circ$, -15°), where perceived tilt was closer to the horizontal (0°).

Since perceived tilt was similar to the projected tilt, it is possible to infer that observers based their judgments primarily on the projected angles in the image plane, and were unable to discount the distortions from the perspective projection in order to recover the veridical tilt. However, under these viewing conditions, projected tilt was a function of internal angle, so it is not possible to discount the possibility that other properties correlated with 3-D shape also played a role in the illusion.

4 Experiment 2

In a second experiment, the viewing conditions were manipulated so that the leading (shorter) edge was perpendicular to the line of sight in the $x-z$ plane ($\omega = 90^\circ$). Projected tilts were therefore constant across stimuli, and independent of internal angle (see figure 7). The leading edges of all the horizontal stimuli ($\psi = 0^\circ$) had the same projections in the image, and these projected tilts were equal to the physical tilts. The projected tilts for nonlevel stimuli are within 2° of the physical tilts. If projected tilt

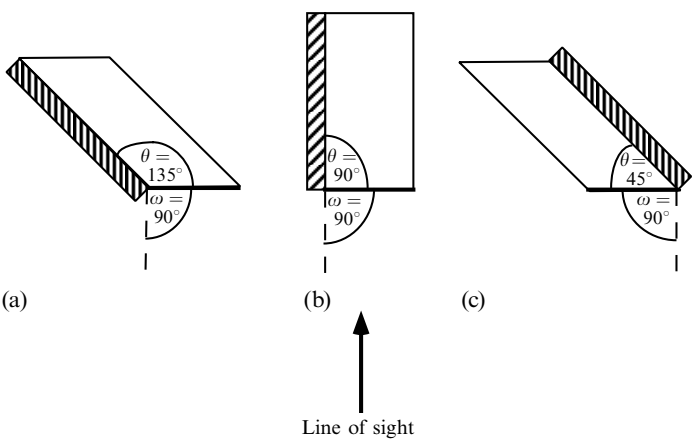

Figure 7. Top view of some of the stimuli presented in the $\omega = 90^\circ$ viewing condition with internal angle, θ, equal to: (a) 135°, (b) 90°, and (c) 45°. Line of sight is perpendicular to the leading edge (thicker lines), with observers facing 'due North', towards the top of the page.

was the only factor involved in this illusion, observers in this experiment should have perceived tilt veridically, especially for horizontal stimuli, and perceived tilt should have been constant for shapes with the same physical tilt. Perceived 3-D tilt was estimated by the same procedures as in experiment 1.

Results from this experiment are shown in figure 8. Perceived tilt is plotted as a function of the internal angle, θ, for each of the three physical tilts used: $\psi = 0^\circ$, $+15^\circ$, -15°. Physical tilts are shown as the three horizontal dotted lines. The dashed lines represent the projected tilts of the stimuli with nonzero physical tilts. The projected tilts for the level stimuli are equal to the physical tilt, ie 0°. Perceived tilt was veridical for 90° shapes (rectangles). Despite the fact that projected angles were matched across stimuli, perceived tilt still varied as a monotonically increasing function of the internal angle, θ. Shapes of internal angle sharper than 90° were perceived as tilted down relative to shapes of internal angle wider than 90°. The data form three distinct, shallow S-shaped curves. The curves are steepest in the center and shallow at the extreme values of internal angle. The illusory tilt is not as strong as in the previous experiment. The magnitude of the illusion represents the effects of factors other than the projected tilt. These factors may include projective distortions of other parts of the stimulus, or biases towards perceiving certain shapes, as is discussed in section 7.

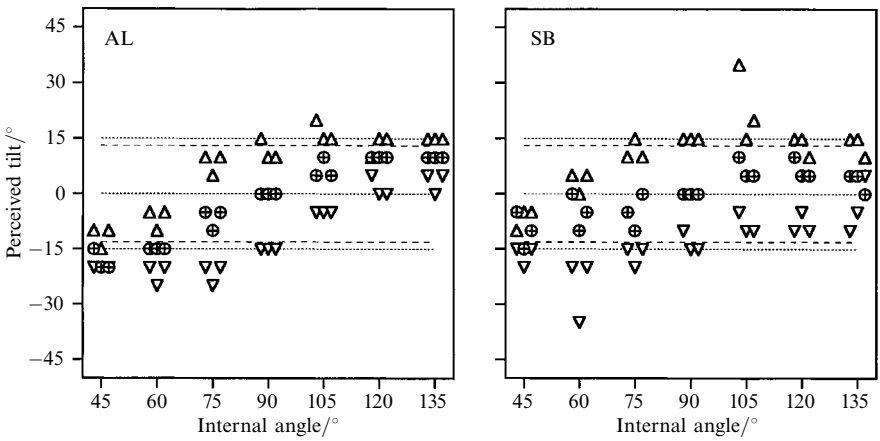

Figure 8. Tilt judgments obtained from experiment 2 ($\omega = 90^\circ$). Results from three physical tilts are presented, including $\psi = 0^\circ$ (no physical tilt), as shown by the crosses-in-circles, $\psi = +15^\circ$ (upwards tilt), as shown by the upward-pointing triangles, and $\psi = -15^\circ$ (downwards tilt), as shown by the downward-pointing triangles. Three measurements are shown for each value of internal angle for each of the three physical tilts. In order to avoid occlusion, two of the three measurements for each internal angle have been shifted horizontally. The three dotted lines represent the physical tilts of the three sets of stimuli. The two dashed lines represent the projected tilts of the two sets of stimuli with physical tilts. The projected tilt of the horizontal ($\psi = 0^\circ$) stimuli is the same as the physical tilt (see text for details).

5 Experiment 3

In the previous experiment, the projections of the leading edges were matched across stimuli, and projected tilts were approximately equal to the physical tilts. Thus, the differences in perceived tilt between the different shapes represent a 'pure' illusion, without any contribution from projective tilt distortions in the projected image.

In the third set of experiments, we varied the experimental conditions so that projections were once again matched across stimuli, but projected tilts were no longer similar to the physical tilts. In effect, we created a conflict between the physical tilt and the projected tilt, as in the first experiment, but this time the size of the discrepancy was constant across all values of internal angle.

Stimuli were presented with the leading edge oriented 45° or 135° relative to the line of sight (in the $x-z$ plane). Figure 9 illustrates parallelogram stimuli as presented in the $\omega = 45°$ viewing condition with internal angles, θ, of 135° (a) and 90° (b). Photographs of actual level stimuli with these internal angles are shown in figures 2a and 2b respectively. The projections of the leading edges in the images are the same, as can be established from measuring the angle between the backboard and the leading edge in both photographs in figure 2.

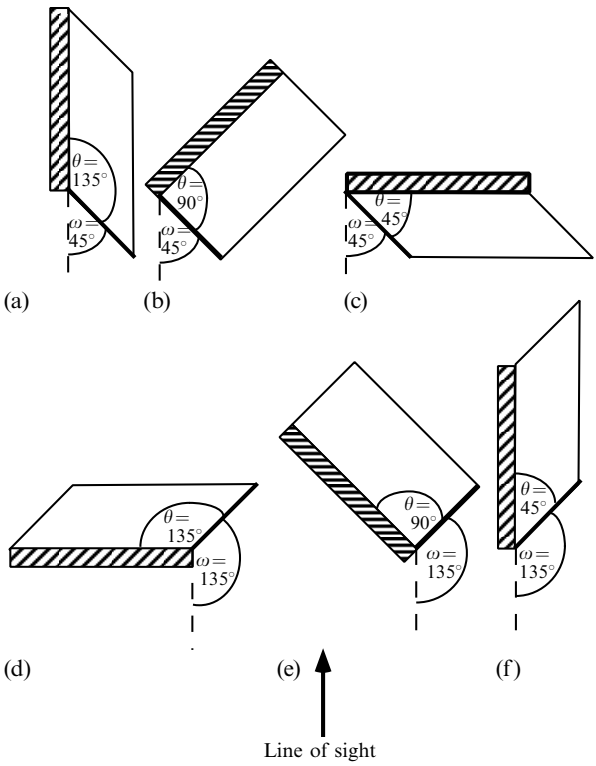

Figure 9. (Top row) Top view of some of the stimuli presented in the $\omega = 45°$ viewing condition, with internal angle, θ, equal to: (a) 135°, (b) 90°, and (c) 45°. Line of sight, as before, is straight up the page. Parts (a) and (b) of the figure represent the orientations of the corresponding stimuli in figure 2. (Bottom row) Top view of some of the stimuli presented in the $\omega = 135°$ viewing condition, with internal angle, θ, equal to: (d) 135°, (e) 90°, and (f) 45°.

Results from experiment 3 are shown in figure 10. Figure 10a shows perceived tilt of each shape as a function of internal angle for the $\omega = 45°$ viewing condition. Perceived tilt is plotted as a function of internal angle for each of the three physical tilts: $\psi = 0°$, $+15°$, $-15°$. Projections of the leading edges were constant across shapes. Projected tilts therefore remained constant, and are represented by the three horizontal dashed lines. Physical tilts are shown by the three horizontal dotted lines. Under these ($\omega = 45°$) viewing conditions, projected tilt did not equal physical tilt (as was the case in the previous experiment), but exceeded it by approximately 25°. Perceived tilt generally varied between the projected tilt and the physical tilt. For shapes with internal angles broader than 90° ($\theta = 105°$ to 135°), perceived tilt was generally closer to the projected tilt, for both observers. For shapes with internal angles sharper than 90° ($\theta = 45°$ to 75°), the magnitude of the illusion was decreased for AL, and perceived tilt remained closer to the physical tilt, compared with the data from the other values of internal angle.

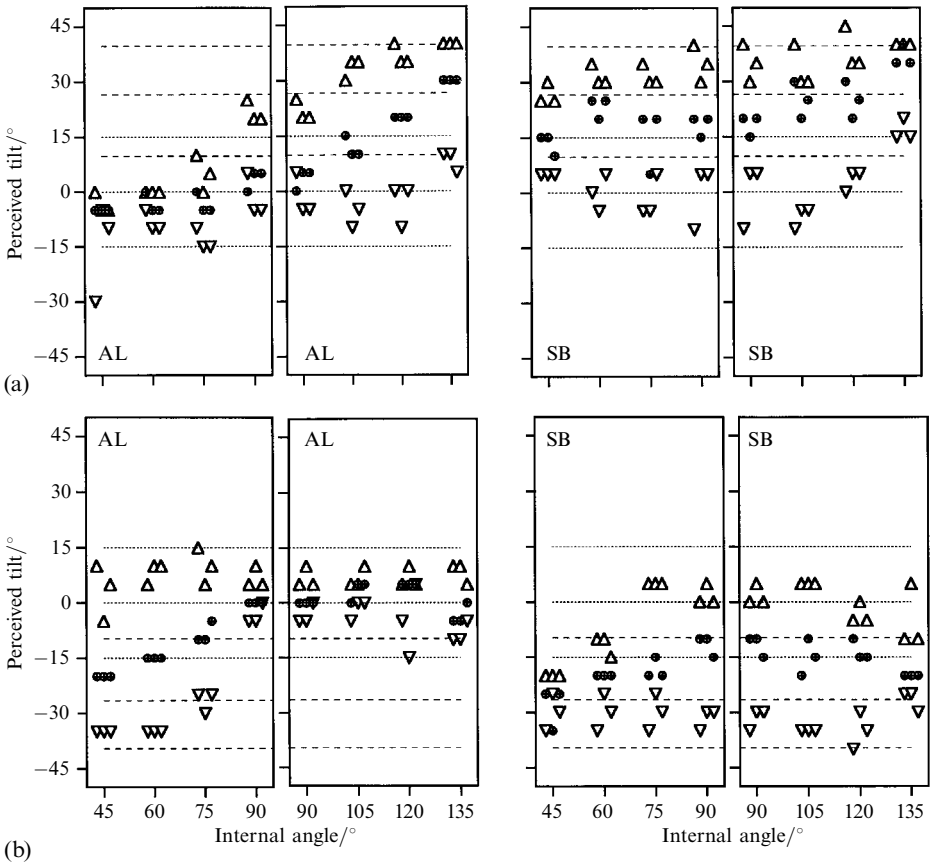

Figure 10. Tilt judgments obtained from experiment 3 for (a) $\omega = 45°$ and (b) $\omega = 135°$. Results from three physical tilts are presented: $\psi = 0°$ (no physical tilt), plotted as crosses-in-circles, $\psi = +15°$ (upwards tilt), plotted as upward-pointing triangles, and $\psi = -15°$ (downwards tilt), plotted as downward-pointing triangles. Three measurements are shown for each value of internal angle for each of the three physical tilts. In order to avoid occlusion, two of the three measurements for each internal angle have been shifted horizontally. The three dashed lines represent the projected tilts, and the dotted lines represent the physical tilts (see text for details).

Observer AL perceived the stimuli with broadest internal angles ($\theta = 135°$) as having a tilt close to the projected tilt, with perceived tilt increasing gradually as internal angle increased, over the range of internal angles greater than 90° (right hand panel). Shapes with internal angles of 90° (rectangles) were perceived as tilted slightly upwards (relative to the physical tilt). Shapes with internal angles of 75° were perceived veridically, while shapes with sharper internal angles were all perceived as having little or no physical tilt (left-hand panel). Observer SB's judgments showed higher values of perceived tilt. Shapes with broader internal angles (greater than 90°) were perceived as having tilts which were similar to the projected tilts (right-hand panel), while shapes with internal angles of 90° or less were seen as having relatively constant tilts in between the projected and the physical tilt (left-hand panel). Under these viewing conditions, the data from both observers are not well described by continuous functions. Instead, it appears as though there are quite different trends in perceived tilt for shapes with internal angles less than 90°, versus internal angles greater than 90°. For this reason, the graphs have been split into two panels, as shown in figure 10a.

Figure 10b shows perceived tilt of each shape as a function of internal angle for the $\omega = 135°$ viewing condition. Perceived tilt is plotted as a function of internal angle

for each of the three physical tilts used: $\psi = 0^\circ$, $+15^\circ$, -15°. Physical tilt is shown by the three horizontal dotted lines. Projected tilt is represented by the three horizontal dashed lines, and remained constant for each physical tilt across all of the shapes used in the experiment. Under these ($\omega = 135^\circ$) viewing conditions, projected tilts were approximately 25° less than the physical tilts. Perceived tilts once again varied between the projected tilt and the physical tilt. For shapes with internal angles 90° or less, perceived tilt more closely resembled projected tilt. For shapes of internal angles greater than 90°, perceived tilt was closer to the physical tilt.

Observer AL perceived all the upward-tilted shapes as having a fairly constant tilt, closer to the physical than the projected tilt. Stimuli with level physical tilt ($\psi = 0^\circ$) were perceived as having a downwards tilt, with this value gradually increasing over the range of values of internal angle up to 90°. Over the range of internal angles from $\theta = 90^\circ$ to 135°, perceived tilt was veridical. Downward-tilted stimuli followed a similar, but more pronounced trend. The separation of points from the different physical tilts was much closer for internal angles of 90° or more. Observer SB showed some similar trends. Perceived tilt for upward-tilted ($\psi = +15^\circ$) stimuli was close to horizontal for most shapes, except for an abrupt decrease of perceived tilt for shapes of internal angle less than 75°. Perceived tilt of downward-tilted ($\psi = -15^\circ$) stimuli remained fairly constant, closer to the projected tilt than the physical tilt. Stimuli with level physical tilt showed a steady increase in perceived tilt for shapes with values of internal angles less than 90°, with a constant, slightly downwards tilt for higher values of internal angle. In general, the functions in figure 10b appear to be 180° rotations of those in figure 10a. In addition, the data once again appear to show different trends, depending on whether internal angle is greater than or less than 90°. Because of this, these graphs have been split into two panels, as in figure 10a.

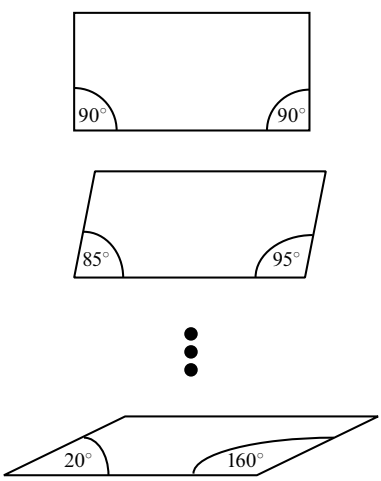

Figure 11. Examples of comparison stimuli used in the shape-matching experiments. From top to bottom, 90° (rectangle), $85^\circ/95^\circ$, and $20^\circ/160^\circ$ parallelograms. (The numbering of the angles is added for the purposes of clarity, and was not present on the comparison stimuli presented to observers.)

The differences between the trends in these two sets of graphs can be attributed, to a certain extent, to the orientation of the stimuli. The $\omega = 45^\circ$ view of the $\theta = 135^\circ$ shape and the $\omega = 135^\circ$ view of the $\theta = 45^\circ$ shape are repetitions of the trials in experiment 1. These views can be termed 'edge' views, since only the edge of the backboard is visible, as illustrated in figures 9a and 9f, respectively. Line of sight for these edge views is parallel to the backboard, and thus parallel to the longer edge of the parallelogram in the $x-z$ plane. Conversely, the $\omega = 45^\circ$ view of the $\theta = 45^\circ$ shape and the $\omega = 135^\circ$ view of the $\theta = 135^\circ$ shape can be termed 'front' or 'back' views, respectively, since only that part of the backboard is visible, as illustrated in figures 9c and 9d, respectively. In the 'front' and 'back' views, the backboard and the longer

edge of the parallelogram were orthogonal to the observer's line of sight in the $x-z$ plane.

Under the $\omega = 45^\circ$ viewing condition, the orientation of the stimulus gradually shifted from 'front' view to 'edge' view, as internal angle increased. The reverse was true for the $\omega = 135^\circ$ viewing condition, in which the orientation of the stimulus shifted from 'edge' view to 'back' view, as internal angle increased. 'Edge' views yielded perceived tilts much closer to the projected tilt. When stimuli were presented in 'front' or 'back' views, observers appeared to correct for projective tilt distortions of the shorter edge, and perceived tilts were much closer to the (veridical) physical tilts. It is also the case that when 'The Future' building (figure 1) is viewed directly from the front, the balconies appear to be level.

6 Experiment 4

After completing all of the tilt-estimation experiments described above, observers performed a shape-estimation experiment. The stimuli and experimental procedure were the same as described in the previous experiments, except that observers were asked to estimate the physical shape of the parallelogram stimuli by selecting the best representative from a series of comparison shapes. There were 14 comparison shapes in all, with internal angles varying, in 5° increments, from a 90° parallelogram (rectangle) to a 20°/160° parallelogram (see figure 11). Observers were allowed to practice the procedure at eye level beforehand, and were also invited to manually examine the stimuli during training.

Observers made shape judgments for each of the three viewing conditions described in experiments 2 and 3 above ($\omega = 90^\circ$, 45°, 135°). Results are shown in figure 12. Perceived shape is plotted as a function of the internal angle, θ, of the shape for the three physical tilts: $\psi = 0^\circ$, $+15^\circ$, -15°. Veridical judgments are represented by the solid lines.

Unlike the data from the previous sets of experiments, no systematic trends can be detected. Observer AL estimated almost all shapes as rectangular. Observer SB's judgments varied randomly between different possible stimulus shapes. In addition, there were no systematic effects of physical tilt, ψ, shown by either observer. In the views presented in these experiments, the projected image was a function of both the shape of the parallelogram, θ, and its tilt, ψ. There is an infinite number of (internal angle, physical tilt) pairs which all give rise to similar projections of their leading edges, ie to similar projected tilts. It seems that our observers had difficulty identifying the veridical physical shape from among other such shapes.

7 Models of the tilt illusion

As shown by the perspective projections derived in Appendix 1, the leading edges of physically level parallelograms are, in general, projected at an angle other than 90° to the projected image of the backboard. Observers are often unable to correct for these kinds of distortions, and this gives rise to an illusory perceived tilt. Since every 2-D projection corresponds to infinitely many 3-D objects, in principle a different shape could be 'seen' for every 'perceived' tilt. The illusory percept is one possible 3-D visual interpretation of the 2-D projection. Another possible interpretation is the veridical physical shape with veridical tilt. The repeatability of the measurements indicates that observers employ a consistent strategy, despite the fact that the set of possible solutions is infinite. In this section we are interested in identifying the specific strategies employed by our observers.

We begin by ruling out a number of simple hypotheses. The illusory tilt of the real balconies (in figure 1), and the results from experiments 1 through 3 serve to rule out the hypothesis that observers were able to extract sufficient information from the retinal

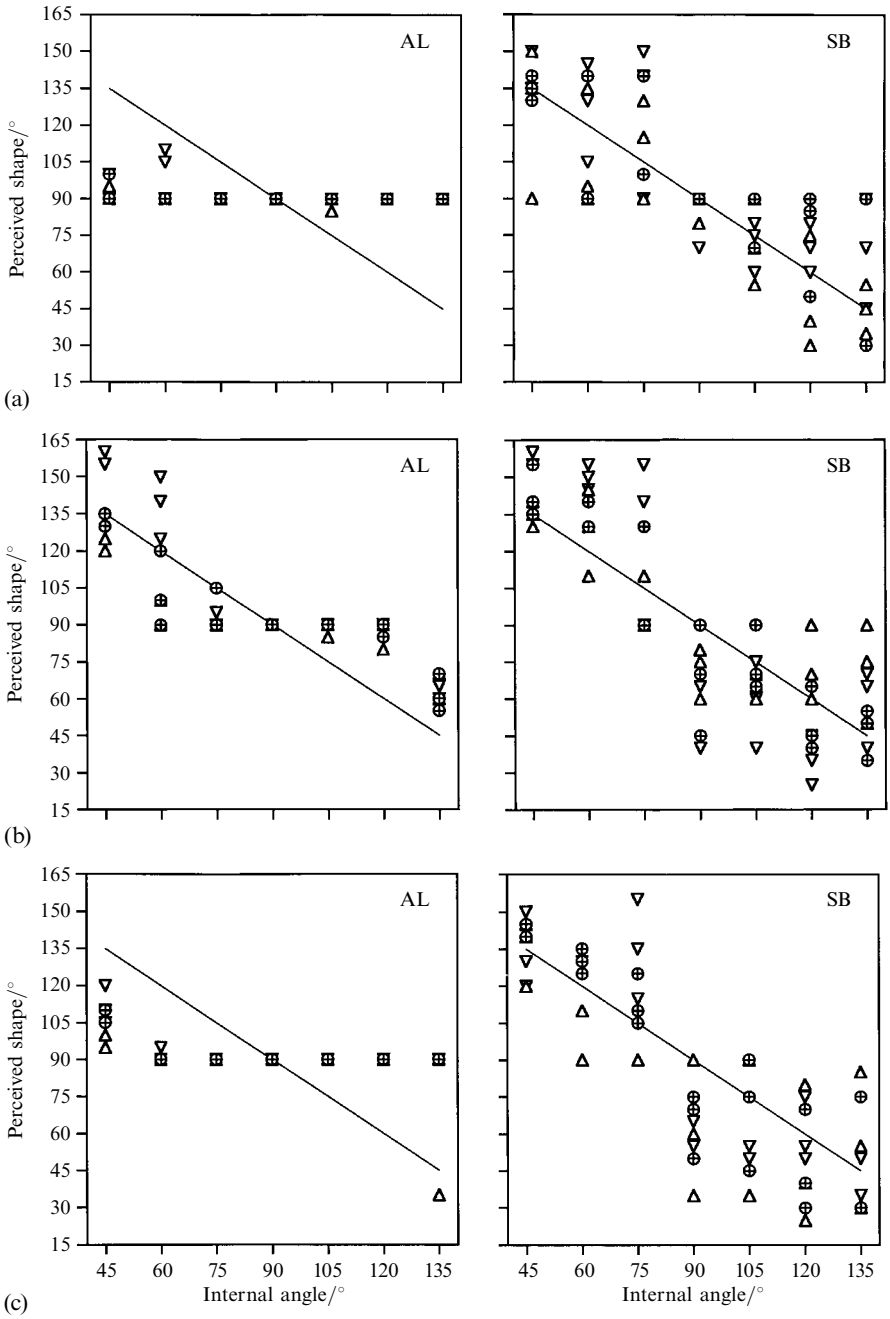

Figure 12. Shape judgments obtained from experiment 4, including (a) $\omega = 90°$, (b) $\omega = 45°$, and (c) $\omega = 135°$ viewing conditions. Results from three physical tilts are presented, including $\psi = 0°$ (no physical tilt), as shown by the crosses-in-circles, $\psi = +15°$ (upwards tilt), as shown by the upward-pointing triangles, and $\psi = -15°$ (downwards tilt), as shown by the downward-pointing triangles. Three measurements are shown for each value of internal angle for each of the three physical tilts. In some cases the symbols are occluded by other points of the same value. Veridical shape judgments are represented by the solid lines.

image to reconstruct the scene veridically. A second hypothesis, that observers could not perceive tilt angles other than those projected in the image, is consistent with the data from experiment 1, but is ruled out by the results from experiment 2, as well as the observation that rectangular balconies still appear level when viewed from multiple viewpoints. In experiment 2, all the leading edges of the stimuli were frontoparallel to the line of sight in the $x-z$ plane, resulting in little or no projective tilt distortions. Despite this fact, illusory tilts were reliably reported by both observers. Experiment 3 showed that neither physical tilt nor projected angles alone can account for the data observed, since perceived tilt varied considerably while these two factors remained constant across different shapes.

Experiments 2 and 3 have shown that, while viewpoint, physical tilt, and projections in the image were held constant, perceived tilt was not constant and varied as a function of internal angle. We therefore conclude that there was a significant effect of stimulus shape in the perception of tilt in these stimuli.

7.1 *Perceptual assumptions*

A recurring theme in visual perception has been the postulation of perceptual tendencies, priors, predispositions, or assumptions to explain percepts that do not correspond to physical reality (Perkins 1972; Richards et al 1996; Mamassian and Landy 1998). For example, the tendency to perceive indeterminately oriented trapezoids as rectangular has been used to explain Ames's room and window demonstrations (Ames 1951). The same tendency would also explain why the slanted trapezoid in figure 2 of Koenderink (1986) appears as a frontoparallel square, and the frontoparallel trapezoid appears as a square slanted in depth. If the observers in our experiments assumed that all the stimulus shapes were rectangular, they would be forced to infer nonveridical tilts to match the projection of the assumed rectangle to that of the physical parallelogram. In its global form, this hypothesis is not plausible for our stimulus conditions because it is impossible for a planar rectangle to share the retinal projection of a nonrectangular parallelogram. The proof of this assertion is presented in Appendix 2.1.

The tendency to see Ames's trapezoids as rectangular is a special case of a spatially local tendency to perceive ambiguously projected 3-D angles as being equal to 90°. Impossible figures like the Penrose triangle (Penrose and Penrose 1958) demonstrate the precedence of local 3-D inferences over global 3-D inferences. In this paper we suggest that observers perceive 2-D projections in terms of consistent 3-D configurations on the basis of local information in the projection and also on the basis of spatially local *perceptual assumptions* that function like default settings.

7.2 *Model 1: Weighted average of rectangularity (v) and level tilt assumptions*

In the first model the observer assumes a priori that, at the junction between the backboard and the leading edge, all the 3-D angles are equal to 90°. In our notation, this corresponds to assuming that $\psi = 0°$, $\theta = 90°$. The implication of this assumption for the perceived shape of the stimulus is that the observer 'sees' a rectangle fitting just the leading edge of the stimulus and the two long edges adjacent to the leading edge. The remaining edge could be fit either by varying the other angles of the parallelogram or by giving up rigidity and bending the planar rectangle along that edge. Informal observations of the perceived shape of the 'balconies' are consistent with either of these possibilities. Appendix 2.2 proves that there is one and only one rectangle (up to parallel displacement) that will share the projections of three sides with a nonrectangular parallelogram. This hypothesized rectangle does not exist in the physical world, but shares the projection of the leading edge with the physical stimulus. We choose to refer to the imagined rectangle as a 'virtual rectangle', and we refer to the leading edge of any such virtual rectangle as a 'virtual edge'. The tilt in world coordinates of the virtual rectangle will generally be different from the physical tilt of the parallelogram. The assumption is

that the observer interprets the projected tilt of the leading edge as arising from the 3-D tilt of the virtual rectangle. The leading edge of the virtual rectangle will be oriented at 90° to the backboard (in the $x-z$ plane), and its virtual tilt, represented by v, is derived in Appendix 3 as a function of internal angle, θ, physical tilt, ψ, elevation, ϕ, and orientation of the leading edge, ω, of the stimulus. The observer thus assumes a particular 3-D shape and infers the 3-D tilt consistent with its projection. However, the same perceptual assumption that local 3-D angles are equal to 90° also implies that the observer assumes that the angle of tilt between the physical object and the backboard is equal to 90°.

In most cases reported here, the predictions based on the shape and tilt implications of the 90° perceptual assumption will be incompatible with one another. To resolve this conflict, we suggest that the observer revises one or both of these predictions in order to infer shape and tilt angles compatible with the projection. The degree of revision of each assumption is determined by the relative uncertainty of information about shape and tilt, which in turn is a function of the viewing angle. For the views presented in our experiments, the amount of shape information which is available for each view of the stimulus is affected by the backboard, which occludes parts of certain stimuli (see figures 7 and 9). Information about the tilt of the leading edge relative to the tilt of the backboard is also dependent on the orientation of the leading edge relative to the observer's line of sight, ω.

The revision process may be quite complex and may involve a number of conditional decisions. However, it has been shown previously (Dawes and Corrigan 1974; Einhorn et al 1979) that, if the correct factors have been identified, the results of complex decision processes can be reasonably approximated by suitable weights in a regression equation. The models in this section will take the form of weighted regression equations, and thus serve as tests of the adequacy of our theory of perceptual assumptions without committing to any particular decision process.

The model is described by equation (2), which states that the perceived tilt, $\hat{\psi}$, of a stimulus shape with internal angle θ and physical tilt ψ, has a value between v, the tilt angle from the vertical of the virtual rectangle, and a 3-D tilt of 90° from the vertical (ie 0°).

$$\hat{\psi}_{\theta,\psi,\omega} = \alpha_\omega v_{\theta,\psi,\omega} + (1-\alpha_\omega)\,0^\circ . \tag{2}$$

The only free parameter in this model, α_ω, is constrained to a value between 0 and 1. The value of α_ω reflects the relative relaxation of assumptions concerning shape and tilt. Higher values of α_ω suggest that the perceived internal angle is closer to 90° and that the perceived tilt corresponds to that of the virtual rectangle. Lower values of α_ω suggest that the shape assumption has been relaxed to a greater extent, and the perceived tilt is closer to 0° (horizontal). The extent to which each assumption is relaxed depends on the amount of information about shape and tilt that the observer derives from each situation. Therefore, we attempted to estimate one value of α_ω for each viewing condition, ie the model was fit separately to the data for each experiment. The values of v were derived from expression (A3.11) (Appendix 3), and the value of α_ω was estimated from the best least-squares fit of the following regression equation:

$$\hat{\psi}_{\theta,\psi,\omega} = \alpha_\omega v_{\theta,\psi,\omega} . \tag{3}$$

Figure 13a shows the fit of the model to the results from experiment 1 ($\omega = 180^\circ - \theta$). Means from the three physical tilts are presented as filled symbols: $\psi = 0^\circ$, $+15^\circ$, -15°. These filled symbols are the means of the corresponding open symbols in figure 6. The dashed curves show the derived values of v and correspond to $\alpha_\omega = 1$ in the model, equivalent to the assumption that the internal angle, θ, was equal to 90° (ie stimulus shapes were rectangular). The horizontal dotted line at 0° tilt from the horizontal corresponds to $\alpha_\omega = 0$ in the model, equivalent to the assumption that the stimuli were

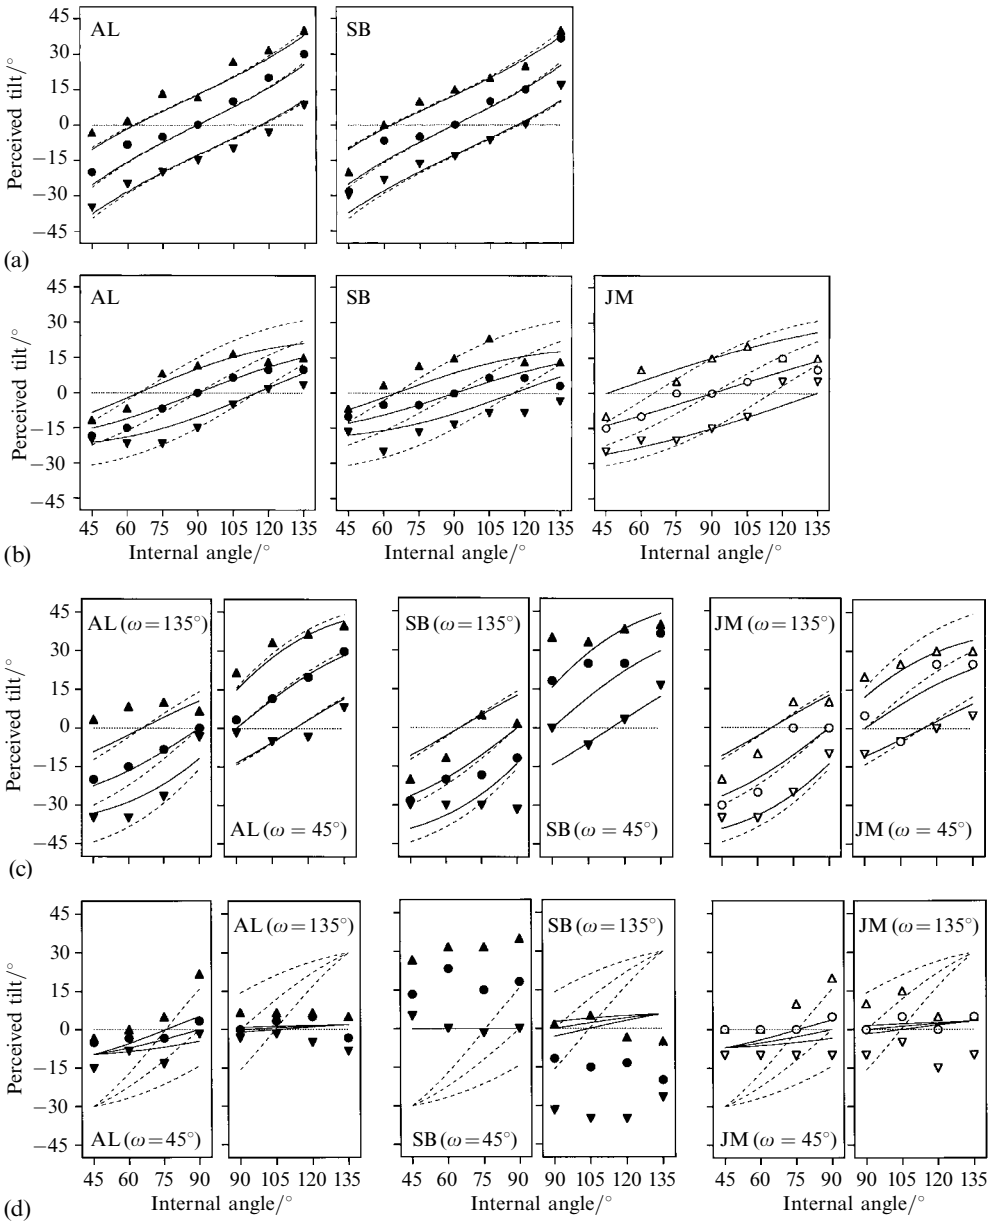

Figure 13. Tilt judgments obtained from different viewing conditions, including (a) $\omega = 180° - \theta$, (b) $\omega = 90°$, (c) and (d) $\omega = 45°$ and $\omega = 135°$. Tilt judgments are plotted as a function of stimulus shape (internal angle). For observers AL and SB, the filled symbols represent means of three sets of observations. In the third panel in (b), (c), and (d), single tilt judgments from observer JM are represented as open symbols. Results from three physical tilts are presented, including $\psi = 0°$ (no physical tilt), as shown by the circles, $\psi = +15°$ (upwards tilt), as shown by the upward-pointing triangles, and $\psi = -15°$ (downwards tilt), as shown by the downward-pointing triangles. Predictions of the best-fitting model (model 1) are shown as three solid curves. The tilts of the virtual rectangle, corresponding to one of the predictions in the model, are shown as three dashed curves, while the 90° prediction is shown as a single dotted line.

level ($\psi = 0^\circ$). For each observer, the value of α_ω was estimated from the best fit of the regression equation to all 21 points in the figure. These best-fitting estimates of α_ω were 0.85 and 0.84 for observers AL and SB respectively. The best-fitting predictions of the model are shown as the solid curves in figure 13a. For both observers, the solid curves appear to provide a good fit to the data, although the fit is only a marginal improvement on the dashed curves representing virtual tilt. It is possible to infer from this model that observers in this experiment seem to be relying almost entirely on the assumption of rectangular shape, and giving very little weight to the assumption of level tilt.

In order to quantify the fit of the model, we calculated average unsigned error of the fit as the average absolute difference in degrees between the best-fitting model and the 21 mean data points. (This measure will be used for all fits because, as will be discussed later, the distribution of errors in the data was too narrow to satisfy the normal-distribution assumptions required by analyses of variance.) The average unsigned errors of the model are 3.72° and 3.69° for AL and SB respectively, indicating that the average error was less than the incremental differences among the comparison stimuli used to estimate the perceived tilts.

Figure 13b shows the model fits obtained from experiment 2 ($\omega = 90^\circ$). The filled symbols represent the mean perceived tilts as in the previous figure, and are the means of the corresponding open symbols in figure 8. The third panel shows the model fits to the data (open symbols) obtained from observer JM. The dashed curves show the derived values of the virtual tilt, v, and the level-tilt assumption is represented by the horizontal dotted line at 0° perceived tilt, corresponding to a (horizontal) tilt of 90° from the vertical backboard. The predictions of the best-fitting model are shown as solid curves. The best-fitting estimates of α_ω were found to be 0.69, 0.58, and 0.69, and average unsigned errors of the model were 2.99°, 4.91°, and 3.68° for observers AL, SB, and JM, respectively. The model provides a good fit to the data obtained from AL and JM, but systematically underestimates the deviations from the horizontal for SB. This is most noticeable in the model predictions for stimuli which have a nonzero physical tilt, represented by the upper and lower solid curves on the graph.

Figures 10a and 10b showed that perceived tilt in experiment 3 was not a simple function of internal angle. In the $\omega = 45^\circ$ viewing condition, perceived tilt increased as a function of internal angle for internal angles greater than 90°, but was relatively constant for smaller angles. The opposite was true for the $\omega = 135^\circ$ viewing condition, in which perceived tilt was relatively constant for angles greater than 90°, but increased as a function of other internal angles. Because of this discontinuity we fit the model to these two ranges ($\theta = 45^\circ$ to 90°, and $\theta = 90^\circ$ to 135°) separately. Figure 13c shows the means of the perceived tilts obtained from the $\omega = 135^\circ$ viewing condition over the range of internal angles $\theta = 45^\circ$ to 90°, alongside means from the $\omega = 45^\circ$ viewing condition over the range of internal angles $\theta = 90^\circ$ to 135°. These views correspond to 'edge' views of the stimuli with 45° and 135° internal angles. Stimuli with other internal angles deviated from this viewpoint by up to a maximum of 45°, as was the case for the rectangular stimuli. Figure 13d shows the means from the $\omega = 45^\circ$ viewing condition over the range of internal angles $\theta = 45^\circ$ to 90°, alongside means from the $\omega = 135^\circ$ viewing condition over the range of internal angles $\theta = 90^\circ$ to 135°. These views correspond to 'front' and 'back' views of the stimuli with 45° and 135° internal angles respectively, with the other stimuli deviating from these viewpoints by up to 45°, as was the case for the rectangular stimuli. In each panel of figures 13c and 13d the dashed curves represent the virtual tilt, v, and the horizontal dotted line represents the assumption of a 90° tilt angle from the backboard (level tilt). The solid curves represent the best-fitting model to all 12 means in that panel. The data from

observer JM (third panel) are shown as open symbols, to indicate that each value represents a single data point.

The best-fitting estimates of α_ω for model fits to data from AL were 0.95 and 0.75 for the $\omega = 135^\circ$ and $\omega = 45^\circ$ viewing conditions, respectively. Average unsigned errors of the model were 3.61° and 4.95°, suggesting a reasonable fit, and the model curves follow the downward trends in the data. Model fits for observer SB are less satisfactory. The best-fitting estimates of α_ω were 1.00 and 0.88. Average unsigned errors for these model fits were 7.38° and 7.12°, about twice those estimated for previous fits. The model fits lie on or close to the gradually decreasing dashed curves representing the virtual tilt, while the perceived tilt remains relatively constant, and exceeds the virtual tilt. Observer JM yielded results similar to AL: best-fitting estimates of α_ω were 0.77 and 0.88 for the two viewing conditions, corresponding to average unsigned errors of 4.38° and 4.40°.

The data shown in figure 13d do not lie close to the dashed curves representing the predictions derived from the tilt of the virtual rectangle. Instead, perceived tilt remained relatively constant across different values of internal angle. The best-fitting estimates of α_ω were much lower than in the previous viewing condition: 0.32 and 0.14 for AL in the $\omega = 45^\circ$ and $\omega = 135^\circ$ viewing conditions, 0.00 and 0.19 for SB, and 0.23 and 0.11 for JM, respectively. For AL and JM, the data fall close to the dotted line representing the horizontal-tilt assumption. This results in the much lower values of α_ω observed. Model fits, as shown by the solid curves, yield average unsigned error values of 3.28° and 5.07° for AL, and 6.32° and 6.91° for JM, respectively. The error values for JM are higher, probably because the model was fitted to only a single set of observations. The model, however, does not fit the data from SB. The model is equivalent to the level-tilt assumption, while the data show significant perceived tilts of opposite sign to those predicted from the virtual tilt. This is reflected in the high values of the error estimates: 16.24° and 16.81° for the $\omega = 45^\circ$ and $\omega = 135^\circ$ data sets, respectively.

In summary, model 1 provides good fits to all of the data from observer AL, with errors ranging from 2.99° to 5.07°. Fits to the data obtained from observer JM had errors ranging from 3.68° to 6.91°. Model fits to the data obtained from observer SB were good for the $\omega = 180^\circ - \theta$ and $\omega = 90^\circ$ viewing conditions, with errors of 3.69° and 4.91°, respectively, but unacceptable for all of the $\omega = 45^\circ$ and $\omega = 135^\circ$ viewing conditions, with errors varying between 7.12° and 16.81°.

7.3 *Model 2: Weighted average of rectangularity (v) and projected tilt (τ) assumptions*
In those cases where model 1 does not fit the observed data, the means of the perceived tilts seem to remain constant across values of internal angle, and lie closer to the projected tilts. In model 2 we retain the perceptual assumption that the shapes were rectangular, but assume that the observers used angles in the projected image in estimating 3-D physical tilt.

The model is described by the equation:

$$\hat{\psi}_{\theta,\psi,\omega} = \alpha_\omega v_{\theta,\psi,\omega} + (1 - \alpha_\omega)\tau_{\psi,\omega}\,, \tag{4}$$

where $\hat{\psi}$ is the perceived tilt, v is the 3-D tilt of the virtual rectangle, and τ is the angle in the image plane between the leading edge and the backboard, ie the projected tilt.

The model fits were generated by finding α_ω, the slope of the best-fitting regression line in equation (4):

$$\hat{\psi}_{\theta,\psi,\omega} - \tau_{\psi,\omega} = \alpha_\omega (v_{\theta,\psi,\omega} - \tau_{\psi,\omega})\,. \tag{5}$$

The best-fitting curves obtained from this model, along with the two predictions from which they were derived, are plotted against the means of the data obtained from experiments 1 through 3 in figure 14.

The filled symbols in figure 14a represent the means of the data from experiment 1 ($\omega = 180^\circ - \theta$). The model in equation (5) results in best estimates of α_ω of 0.00 for both observers, although for these data the parameter α_ω has very little effect on the shape of the curve and the fit of the model. The best-fitting model is shown as the solid curves, which also correspond to the projected tilts. The dashed curves correspond to the virtual-tilt assumptions. It should be noted that the projected tilt, τ, and virtual tilt, v, predictions give rise to very similar curves. Average unsigned errors for this model are 3.56° for AL and 3.71° for SB. These values are similar to those obtained from model 1. In other words, whereas experiment 1 clearly documented the illusion, the results cannot distinguish between a reliance on a rectangularity assumption and a reliance on the projected tilt.

In figure 14b the filled symbols show the means of the data obtained from experiment 2 ($\omega = 90^\circ$). The data from MA (third panel) are shown as open symbols, to indicate that each value represents a single data point. The virtual tilts, v, are shown as the three sets of dashed curves in this and the following two figures. The projected tilts, τ, are shown as the sets of horizontal dotted lines. The best fits of model 2 are shown as solid curves. Model 2 provides a good fit to the data obtained from the $\omega = 90^\circ$ view, although for AL perceived tilt is overestimated for the more extreme values of internal angle ($\theta = 45^\circ$ to 60°, 120° to 135°). The best estimates of α_ω were 0.68 for AL, 0.41 for SB, and 0.63 for MA. These values are similar to those obtained with model 1, as are the average unsigned errors of the model of 3.51°, 2.96°, and 4.69° for AL, SB, and MA, respectively. Under these viewing conditions, model 2 provides a slightly better fit to the data obtained from SB and MA, while AL's data are better explained by model 1.

The differences in prediction of the two models are most evident in the results of experiment 3. The means from AL and SB are shown as the filled symbols in figures 14c and 14d. The data from MA (third panel) are shown as open symbols, to indicate that each value represents a single data point. In figure 14c, the model (represented by the solid curves) provides good fits to the data for five out of the six panels. The exception is AL's data from the $\omega = 135^\circ$ viewing condition (first panel), where the predictions from model 2 consistently underestimate the perceived tilt. This fit gives rise to the average unsigned error value of 7.33°, higher than that obtained from model 1. For the $\omega = 45^\circ$ viewing condition, the α_ω value of 0.78 for AL is less than that estimated in model 1. This finding suggests that the level-tilt prediction is less useful in predicting the perceived tilt, and a better fit is obtained from the projected-tilt and virtual-tilt assumptions. The average unsigned error of 3.29° is slightly lower than that obtained from model 1. Model 2 provides a better fit than model 1 to the data from SB and MA under these viewing conditions. Best estimates of α_ω were 0.40 and 0.53 for SB, and 0.65 and 0.72 for MA, for the $\omega = 135^\circ$ and $\omega = 45^\circ$ viewing conditions, respectively. Average unsigned errors of the model were 3.76° and 4.79° for SB, and 4.09° and 3.51° for MA. These values are much lower than those obtained from model 1.

In figure 14d, the model 2 fits, as represented by the solid curves, are very different from the corresponding model 1 fits shown in figure 13d. From the first two panels one can see that the data points for observer AL remain relatively constant at around 0° of perceived tilt for all three physical tilts, but the best-fitting curves increase as a function of internal angle. The fit to AL's data from the $\omega = 135^\circ$ viewing condition, corresponding to a best estimate of α_ω equal to 0.69, is much worse than model 1, and yields an average unsigned error of 11.72°. The fit to the data from the $\omega = 45^\circ$ viewing condition is somewhat better than model 1, and the best estimate of α_ω of 0.65 yields an average unsigned error of 3.99°.

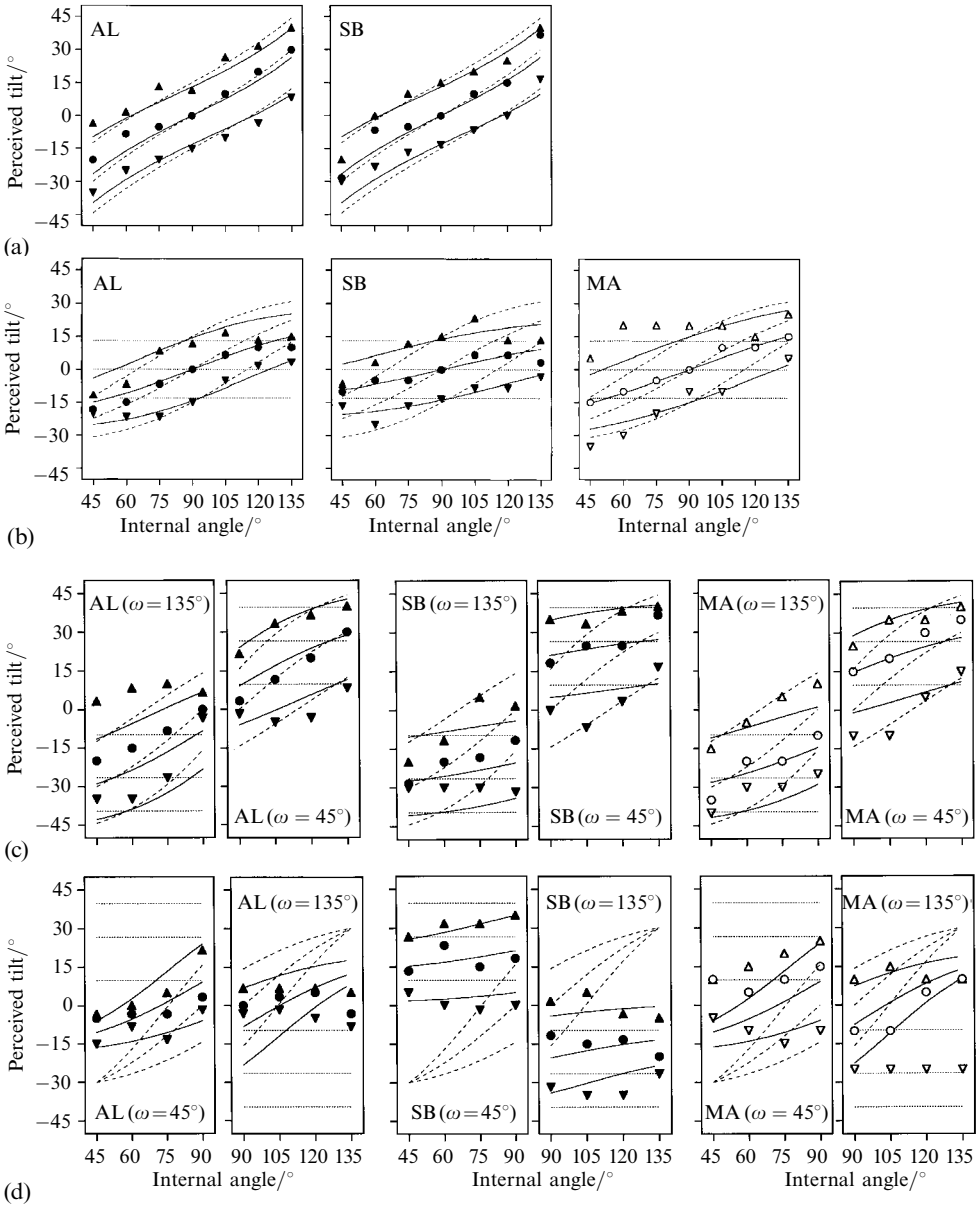

Figure 14. Tilt judgments obtained from different viewing conditions, including (a) $\omega = 180° - \theta$, (b) $\omega = 90°$, (c) and (d) $\omega = 45°$ and $\omega = 135°$. Tilt judgments are plotted as a function of stimulus shape (internal angle). For observers AL and SB, the filled symbols represent means of three sets of observations. In the third panel in (b), (c), and (d), single tilt judgments from observer MA are represented as open symbols. Results from three physical tilts are presented, including $\psi = 0°$ (no physical tilt), as shown by the circles, $\psi = +15°$ (upwards tilt), as shown by the upward-pointing triangles, and $\psi = -15°$ (downwards tilt), as shown by the downward-pointing triangles. Predictions of the best-fitting model (model 2) are shown as three solid curves. The tilts of the virtual rectangle, corresponding to one of the predictions in the model, are shown as three dashed curves, while the projected tilts are shown as dotted lines. In figure 14a, the predictions of the model are equivalent to the projected tilt, and the two sets of curves are superimposed.

For observers SB and MA, however, the model 2 fits show a marked improvement over model 1. The solid curves in the figure represent the best fit of the model. These curves show only a slight increase in predicting tilt for increasing values of internal angle, and correspond to the trends in the data, for both SB and MA. The best-fitting values of α_ω were 0.23 and 0.20 for SB, and 0.45 and 0.45 for MA, for the $\omega = 45^\circ$ and $\omega = 135^\circ$ viewing conditions, respectively. The model fits resulted in average unsigned errors of 5.99° and 2.97° for SB, and these errors are lower than those obtained from model 1. Errors for MA were 3.70° and 11.14°, for the $\omega = 45^\circ$ and $\omega = 135^\circ$ viewing conditions, respectively. While these values are much lower than those obtained from model 1, MA's data from the $\omega = 135^\circ$ viewing condition are noisy, and cannot be fit adequately by either model.

7.4 *Comparing the models*

In order to facilitate a comparison of the two models, the best-fitting values of α_ω for each of the viewing conditions are presented in table 2, along with the average unsigned errors from each model. The $\omega = 45^\circ$ and $\omega = 135^\circ$ orientation fits are split into two columns: the first two values correspond to the model fits shown in figures 13c and 14c ($\omega = 45^\circ$, $\theta = 90^\circ$ to 135°; $\omega = 135^\circ$, $\theta = 45^\circ$ to 90°); the second two values correspond to the model fits shown in figures 13d and 14d ($\omega = 45^\circ$, $\theta = 45^\circ$ to 90°; $\omega = 135^\circ$, $\theta = 90^\circ$ to 135°).

The average of the absolute differences between the predictions and the data was used as a measure of fit, because the conditions for the more traditional analysis of variance were not met. The within-subjects error in the experiment was very low, even though trials from experiments 1 through 3 were interleaved and presented in random order. The errors were not normally distributed, as is required for an analysis of variance. The error distributions from observers AL and MA each had over three times the kurtosis of a normal distribution with the same mean and standard deviation as the data set. This may have been due to the coarse spacing of the angles in the set of comparison stimuli. Since the comparison stimuli were in 5° tilt increments, we considered any model fit with an error less than 5° as reasonable for these observers.

Applying this measure, table 2 indicates which model fits are acceptable for each of the viewing conditions presented to observers AL and SB. Both models provide acceptable fits for AL with no clear differences between errors, except for the model 2 fits to the means of the data from the $\omega = 135^\circ$ viewing conditions, for which the errors are two to four times the size of the errors obtained from model 1. For this reason alone, we conclude that model 1 is better for AL's data. For observer SB's data, model 2 consistently shows lower errors. There is no difference between the two models for the data from experiment 1, and only the results from experiments 1 ($\omega = 180^\circ - \theta$) and 2 ($\omega = 90^\circ$) are reasonably fit by model 1. We therefore conclude that model 2 is better for SB's data. For JM and MA, we found a similar pattern in that both models

Table 2. Values of α (shown in bold) and average unsigned error (shown in italics) obtained from the best fits of the two models to the data for observers AL and SB.

View (ω)	$\omega = 180^\circ - \theta$	$\omega = 90^\circ$	$\omega = 45^\circ$	$\omega = 135^\circ$	$\omega = 45^\circ$	$\omega = 135^\circ$
Shape (θ)	(45° to 135°)	(45° to 135°)	(90° to 135°)	(45° to 90°)	(45° to 90°)	(90° to 135°)
Observer AL						
model 1	**0.85** *3.72*	**0.69** *2.99*	**0.95** *3.61*	**0.75** *4.95*	**0.32** *5.07*	**0.14** *3.28*
model 2	**0.00** *3.56*	**0.68** *3.51*	**0.78** *3.29*	**1.00** *7.33*	**0.65** *3.99*	**0.69** *11.72*
Observer SB						
model 1	**0.84** *3.69*	**0.58** *4.91*	**1.00** *7.38*	**0.88** *7.12*	**0.00** *16.81*	**0.19** *16.24*
model 2	**0.00** *3.71*	**0.41** *2.96*	**0.40** *3.76*	**0.53** *4.79*	**0.20** *2.97*	**0.23** *5.99*

provided reasonable fits to the data from experiment 2 ($\omega = 90^\circ$ viewing condition), while for experiment 3, model 1 provided better fits for JM and model 2 provided better fits for MA.

In both of the models, an α_ω value of 0.5 represents equal weighting to both predictions. Higher values of α_ω suggest that more weight is given to the tilt of the virtual rectangle than the tilt assumption (90° in model 1, projected tilt in model 2). In general, α_ω values are higher than 0.5. This is especially true for model 1 in those cases where average unsigned error is lowest. We take this finding to imply that, for all observers, *shape assumptions play the larger role in the perceived shape of 3-D objects*. The stimuli we have constructed present the visual system with a conflict between the information contained in the image on the retina, projected from the 3-D real world, and what is assumed. In most cases, the assumption wins, and the result is the illusion observed.

Modeling the data from experiment 1 ($\omega = 180^\circ - \theta$) yields high α_ω values. In this viewing condition, since projected tilt was clearly not level, and internal angles were less visible, there was little reason to revise the shape assumption, thus giving rise to the illusory tilt. Concerning the other viewing conditions, $\omega = 90^\circ$ data yield median values of α_ω for each observer, the $\omega = 45^\circ$ (internal angles of 90° to 135°) and $\omega = 135^\circ$ (internal angles of 45° to 90°) viewing conditions yield the highest values of α_ω, while the remaining data from the $\omega = 45^\circ$ and $\omega = 135^\circ$ viewing conditions yield the lowest value of α_ω. The magnitude of these values of α_ω is related to the amount of the stimulus shape which was visible under the different viewing conditions. Specifically, those views which can be described as 'front' or 'back' views of the stimulus lead to much lower values of α_ω, suggesting a preference for the alternative assumption, whereas 'edge' views (experiment 1) contain the least information about the shape of the stimulus. In this sense, the value of α_ω is an inverted U-function of the angle between the observer's line of sight and the backboard, with its peak at 0 (experiment 1, 'edge' views), and its minima at those stimulus/viewpoint combinations previously identified as 'front' or 'back' views, in which the backboard is orthogonal to the line of sight. In general, the more of the parallelogram shape that is seen, or the more evidence there is that it is nonrectangular, the more likely an observer is to modify the rectangularity assumption. This implies that more weight will be given to the other assumptions, resulting in lower values of α_ω. This holds true for both models, and all observers.

8 Discussion

We have examined a shape illusion, in which the balconies of a building appear to tilt up or down, depending on the viewpoint. The balconies are actually level nonrectangular parallelogram shapes, but appear as tilted rectangles. This effect was replicated in the laboratory under controlled conditions, and it was shown that illusory tilt is a function of both the orientation and the physical shape of the stimulus.

Distortions occur under perspective projection, and these can lead to illusory tilt for nonfrontoparallel edges. However, in experiment 2, we have presented a set of stimuli which give rise to illusory tilts even when the leading edges are adjusted to be in the frontoparallel plane. The illusory tilts observed in experiment 3 also cannot be due to projective distortions, since the orientations of the leading edges were held constant, leading to a constant projected tilt, whereas perceived tilt varied as a function of stimulus shape.

The use of parallelogram stimuli in this study provides a more stringent test of perceptual assumptions than the use of Ames-type trapezoids. First, unlike a trapezoid, a planar rectangle cannot share projections with a nonrectangular parallelogram. The use of a rectangularity assumption by our observers cannot therefore be attributed to 'good form', as in the explanation provided by Perkins (1976), since the rectangle can be matched to only three sides of the stimulus shape, and the fourth side must either

not be at right angles to the adjacent sides or must be nonplanar. In addition, by varying viewpoint, we were able to estimate the degree to which the local rectangularity assumption was used by our observers.

We have introduced a model in which the observer assumes that the leading edge of the stimulus is part of a rectangle with a tilt that yields the same projected tilt of the leading edge as the stimulus shape itself. This virtual tilt correlated highly with the perceived tilt. This assumption, however, was not sufficient to model the data, and the two models also incorporate a priori assumptions about the 3-D tilt of the stimulus shapes.

The fact that the model was successful for all viewing conditions, but with a different value of α_ω, implies that our observers used a different, yet consistent, strategy under each viewing angle, despite all the viewing angles being randomly interleaved in the presentation. Values of α_ω in the model, corresponding to the relative weighting of the shape assumption, varied as a systematic function of viewing angle. The highest values of α_ω were obtained when most of the stimulus shape was occluded by the backboard, and the least amount of shape information was available to the observer. We infer from this that the hypothesis that observers perceive ambiguous parallelogram shapes as rectangular is valid. The lower α_ω values for observer SB (and, to a lesser extent, observer MA) are also consistent with the results of experiment 4. In this experiment, SB showed much greater variability in shape judgments. A less consistent perceived shape might result in less weight being applied to the assumption that the parallelogram stimuli were rectangular.

The results of this experiment suggest that the rectangularity assumption is pervasive, and observers are willing to sacrifice level tilt from the horizontal, despite cognitive implausibility, in order to satisfy this assumption. One point worth making is that this illusion also works well when stimuli are tilted 90° and viewed from one side, so that the backboard is horizontal, and the illusion is of tilt from the vertical. It should also be emphasized that while the observers in these experiments were restricted in their view of the stimulus, forcing them to make a single perceptual judgment, this illusion is persistent in natural settings. Even if one views the building in figure 1 from several viewpoints and is completely familiar with the nature of the illusion, the illusion still remains, and the veridical percept still eludes even the most experienced observers.

To explain the data, we used the notion of perceptual assumptions about the nature of objects in a given display. Our analysis suggests that these prior assumptions are fixed within the visual system, and have not arisen in response to the demands of a particular task. Before the experiments were conducted and the data collected, our observers were shown the full range of stimuli to be used in the experiments. However, the observers seemed to use rectangles rather than the complete set of shapes as the prior assumptions in making tilt judgments. In the shape estimates (experiment 4), at least one of the observers spanned the range of shapes initially presented. It is possible that cognitive priors can be influenced easily by exposure to a range of stimuli, but that the priors used by the visual system in shape inferences are much less malleable. It has been pointed out that the projected image of the stimulus is compatible with an infinite number of possible stimulus shapes, and our model incorporates a weighting function which may be taken to imply a tradeoff between perceived shape and tilt of the stimulus. This possibility is not supported by the data from experiment 4, however, where there is no systematic relationship between perceived shape and perceived tilt (as measured by experiments 2 and 3). There are a number of possible reasons for this. Despite using the best procedure we could, we were unable to obtain clear measurements of perceived shape, and the shape estimation data are insufficient to test any particular hypothesis. It seems to us that it is almost impossible to obtain good estimates of perceived global 3-D shapes.

We have demonstrated that observers have a tendency to 'see' ambiguously projected quadrilaterals as rectangular, with internal angles of 90°. What is not entirely clear from these experiments, however, is whether this tendency is dependent on the number of edges or vertices of a given shape, or whether it is limited to quadrilaterals. Additional informal observations suggest that illusions such as this occur over a wide range of shapes, such as triangles, and are not even dependent on edges being straight. Furthermore, if one moves illusory stimuli such as these up and down in the observer's field of view, the illusion is not diminished. Rather, the percept is one of nonrigid motion, in which the stimulus appears to change shape as it moves (Griffiths and Zaidi 1998b).

In conclusion, very little shape information is invariant under perspective projection. Straight lines project as straight lines and cross-ratios are preserved, but stimulus properties such as symmetry, orthogonality, and parallelism are all lost. Therefore, it is remarkable that human observers are able to infer veridical 3-D shapes most of the time. We have identified stimuli for which veridical inferences are made, and others which lead to a false percept. On the basis of our observations we present a theory that the human visual system uses prior assumptions in order to infer shape. Where these assumptions correspond to the test stimulus, the percept is veridical. Where these assumptions are in conflict with properties of the stimulus, perceptions can deviate greatly from the veridical, as evidenced by the resulting illusions.

Acknowledgements. This work was supported in part by Schnurmacher Institute for Vision Research grants #96-97-075 and #97-98-093, awarded to Q Zaidi. A portion of this work was reported at the ARVO Annual Meetings (Griffiths and Zaidi 1997, 1998a), and forms part of the first author's PhD dissertation. We are grateful to Henry Pinkham of Columbia University for outlining the proofs described in Appendix 2, and to Thomas Papathomas, Hal Sedgwick, Barbara Gillam, Jacob Feldman, Michael Leyton, and members of a seminar group for extended discussions, and to an anonymous reviewer for multiple suggested improvements that were incorporated into the final text. We are also grateful to Andrea Li, Sumati Buradagunta, Molly Aitken, and Jeanette Meng for acting as observers.

References

Ames A Jr, 1951 "Visual perception and the rotating trapezoidal window" *Psychological Monographs* **65**(7) whole number 324

Dawes R M, Corrigan B, 1974 "Linear models in decision making" *Psychological Bulletin* **81** 95 – 106

Einhorn H J, Kleinmutz D N, Kleinmutz H, 1979 "Linear regression and process tracing models of clinical judgment" *Psychological Review* **86** 465 – 485

Griffiths A F, Zaidi Q, 1997 "Three-dimensional illusions: Projections and priors" *Investigative Ophthalmology & Visual Science* **38**(4) S1001e

Griffiths A F, Zaidi Q, 1998a "A model of 3-dimensional shape and tilt perception" *Investigative Ophthalmology & Visual Science* **39**(4) S854d

Griffiths A F, Zaidi Q, 1998b "Rigid objects that appear to bend" *Perception* **27** 799 – 802

Halper F, 1997 "The illusion of The Future" *Perception* **26** 1321 – 1322

Ittelson W H, 1952 *The Ames Demonstrations in Perception* (Princeton, NJ: Princeton University Press)

Koenderink J J, 1986 "Optic flow" *Vision Research* **26** 161 – 180

Koffka K, 1963 *Principles of Gestalt Psychology* (New York: Harcourt, Brace & World)

Mamassian P, Landy M S, 1998 "Observer biases in the 3d interpretation of line drawings" *Vision Research* **38** 2817 – 2832

Mundy J L, Zisserman A, 1992 "Projective geometry for machine vision", in *Geometric Invariance in Computer Vision* Eds J L Mundy, A Zisserman (Cambridge, MA: MIT Press) pp 463 – 519

Penrose L S, Penrose R, 1958 "Impossible objects: A special type of visual illusion" *British Journal of Psychology* **49** 31 – 33

Perkins D N, 1972 "Visual discrimination between rectangular and nonrectangular parallelepipeds" *Perception & Psychophysics* **12** 396 – 400

Perkins D N, 1976 "How good a bet is good form?" *Perception* **5** 393 – 406

Richards W, Jepson A, Feldman J, 1996 "Priors, preferences and categorical percepts", in *Perception as Bayesian Inference* Eds D C Knill, W Richards (Cambridge: Cambridge University Press) pp 80 – 111

Sedgwick H A, Nicholls A L, 1993 "Cross talk between the picture surface and the pictured scene: Effects on perceived shape" *Perception* **22** Supplement, 109

Sedgwick H A, Nicholls A L, Brehaut J, 1995 "Perceptual interaction of surface and depth in optically minified pictures" *Investigative Ophthalmology & Visual Science* **36**(4) S667

Appendix 1. Projected tilt of the leading edge

In all the derivations to follow, the general situation is depicted in figure A1.1. In 3-D 'world coordinates', x denotes lateral (left – right) position, y denotes height, and z depth, relative to the observer's eye. The focal point is at (0, 0, 0). The leading (shorter) edge of each parallelogram is denoted by AB, where the point A corresponds to the intersection of the leading edge with the backboard, and point B is the opposite end of the leading edge. The coordinates of A and B are (x_A, y_A, z_A) and (x_B, y_B, z_B). Without loss of generality, we can assume that $x_A = 0$, ie A is directly in the line of sight, and that the length of AB is one unit. In order to determine the projected tilt of line AB, we determine the slope of the line A′B′, where A′ and B′ correspond to the projections of A and B onto the image plane. The image plane is orthogonal to the line of sight, OA, at a distance of one unit from the focal point.

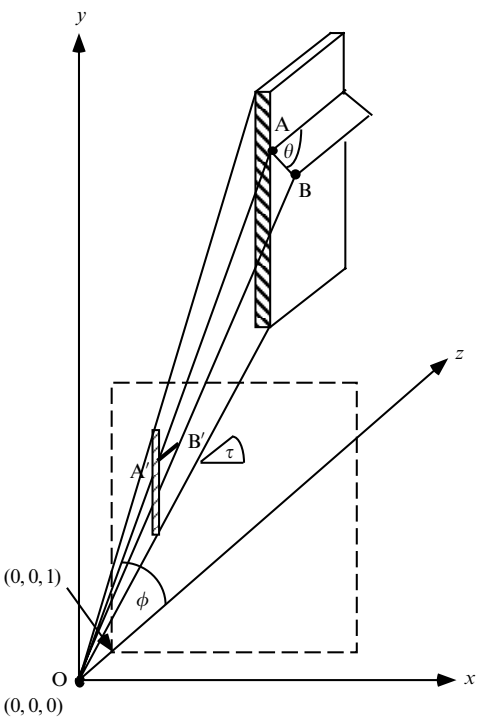

Figure A1.1. Calculation of the projected tilt. In our world coordinates, x denotes lateral (left – right) position, y denotes height, and z depth, relative to the observer's eye, with the focal point O at (0, 0, 0). The stimulus, shown here in the $180^\circ - \theta$ configuration with internal angle, θ, equal to 135° (therefore, $\omega = 45^\circ$) and level ($\psi = 0^\circ$) tilt, is projected onto the $x - y$ image plane, at $z = 1$. The leading (shorter) edge of the parallelogram is denoted by AB, while its projection in the image plane is denoted by the line A′B′. The angle of elevation, ϕ, represents the angle between OA and the z-axis. The projected tilt is represented by τ, the slope of the line A′B′ in the image plane.

The world coordinates of the two points A and B of a leading edge presented at an angle ω to the line of sight are given by:

$$A = (0,\ z_A \tan\phi,\ z_A) \qquad B = (\sin\omega\cos\psi,\ z_A \tan\phi + \sin\psi,\ z_A - \cos\omega\cos\psi). \quad \text{(A1.1)}$$

In the first experiment, ω, the angle between the line AB and the line $x = 0$ in the horizontal plane, was equal to $180^\circ - \theta$. In the other experiments, the value of ω was constant. The physical tilt of the edge, the angle between AB and the horizontal ($y = y_A$) plane, is represented by ψ. ϕ represents the angle of elevation of line of sight from the observer's eye to A; therefore $\tan\phi = y_A / z_A$.

Since the line of sight is elevated, rather than horizontal, the image plane is not vertical, but tilted back by an angle equal to the angle of elevation, ϕ. To calculate the projection onto the nonvertical image plane, we simplified the derivation by rotating the y- and z-axes about the x-axis, by an angle equivalent to ϕ. This has the effect of rotating the image plane onto a vertical image plane at $z = 1$.

Under the rotational transformation of the axes just described, each point P, with coordinates (x_P, y_P, z_P) is mapped onto the point P*, where:

$$\begin{aligned} x_P^* &= x_P \,, \\ y_P^* &= y_P \cos\phi - z_P \sin\phi \,, \\ z_P^* &= y_P \sin\phi + z_P \cos\phi \,. \end{aligned} \tag{A1.2}$$

If (A1.1) is substituted into (A1.2), the coordinates of A and B are now given by:

$$\begin{aligned} A^* &= (0,\ 0,\ z_A \cos\phi + z_A \tan\phi \sin\phi) \\ B^* &= \begin{pmatrix} \sin\omega\cos\psi,\ \sin\psi\cos\phi + \sin\phi\cos\omega\cos\psi \,, \\ z_A\cos\phi - \cos\omega\cos\psi\cos\phi + z_A\tan\phi\sin\phi + \sin\psi\sin\phi \end{pmatrix}. \end{aligned} \tag{A1.3}$$

The projections A′ and B′, correspond to the image coordinates $(X_{A'}, Y_{A'})$ and $(X_{B'}, Y_{B'})$, where:

$$(X_{A'},\ Y_{A'}) = \left(\frac{x_A^*}{z_A^*},\ \frac{y_A^*}{z_A^*}\right), \quad (X_{B'},\ Y_{B'}) = \left(\frac{x_B^*}{z_B^*},\ \frac{y_B^*}{z_B^*}\right). \tag{A1.4}$$

In the image plane, the projection of the backboard is vertical, so the angle between A′B′ and the projection of the backboard is equal to 90° minus the slope of A′B′. The slope of A′B′ is equal to

$$\frac{Y_{B'} - Y_{A'}}{X_{B'} - X_{A'}} = \tan\tau \,, \tag{A1.5}$$

ie τ is the 'projected tilt' of AB, relative to the horizontal.

Since point A′ now lies on the z-axis, equation (A1.5) simplifies to

$$\tan\tau = \frac{y_B^*}{x_B^*} \,. \tag{A1.6}$$

Substituting (A1.3) through (A1.4) into (A1.6), and simplifying, we get

$$\tau = \arctan\left(\frac{\tan\psi\cos\phi + \cos\omega\sin\phi}{\sin\omega}\right). \tag{A1.7}$$

Appendix 2. Proofs of two assertions

A2.1. To every quadrilateral projection corresponds one and only one parallelogram shape.

This assertion implies that the retinal projection of a planar parallelogram can not also be the projection of any other planar parallelogram with different internal angles. In particular, a planar, nonrectangular parallelogram and a planar rectangle can not share the same retinal projection. Consequently, it is geometrically impossible for an observer to 'see' the projection of a planar parallelogram as belonging to a planar rectangle.

For a monocular observer, assume a focal point O and a planar retina on which is projected a planar quadrilateral object with vertices P_1, P_2, P_3, P_4. The observer thus can only see that the vertices of the projected object lie on the rays R_i passing through O and P. The conjecture is that there are planes H in object space such that the intersections Q_i of H with the rays R_i form a parallelogram, and that H is uniquely determined up to a parallel displacement in distance from the observer. In other words, even though a particular quadrilateral retinal projection can arise from infinitely many planar quadrilateral objects, this set of quadrilaterals includes one and only one class of parallelograms with the same internal angles.

Proof: Let K_{12} be the unique plane containing R_1 and R_2; K_{23}, K_{34}, and K_{41} defined analogously. Let L be the line of intersection of K_{12} and K_{34}, and M the line of intersection of K_{23} and K_{41} (the ray-planes defining the 'opposite' sides of the quadrilateral).

Suppose we have found a parallelogram $Q_1Q_2Q_3Q_4$. The line Q_1Q_2 contained in the plane K_{12} is parallel to the line Q_3Q_4 contained in K_{34}, so they meet 'at infinity'. In particular they do not meet the line L of intersection of K_{12} and K_{34}. Therefore, Q_1Q_2 is parallel to L in the plane K_{12} and Q_3Q_4 is parallel to L in the plane K_{34}. In the same way, Q_2Q_3 is parallel to M in the plane K_{23} and Q_4Q_1 is parallel to M in the plane K_{41}. This makes the construction of $Q_1Q_2Q_3Q_4$ unique up to parallel displacement in the plane formed by L and M. It is clear that there is one and only one class of planar parallelograms for every retinal projection that is a quadrilateral. QED.

A2.2. One and only one planar rectangle can project to three sides of the projection of a nonrectangular parallelogram.

If we take the same setup as the previous proof, we now attempt to slice the rays R_i by a plane H so that the quadrilateral $Q_1Q_2Q_3Q_4$ has at least one pair of parallel sides, and a fixed one of the remaining two sides (the 'leading edge') perpendicular to the two parallel sides. Then again there is a unique plane H (up to parallel displacement) that makes this happen. The proof is similar.

Proof: Two sides of the quadrilateral $Q_1Q_2Q_3Q_4$ are constrained to be parallel to the intersection line of the ray-planes formed by the corresponding opposite sides of the retinal projections $P_1P_2P_3P_4$, and the third side of $Q_1Q_2Q_3Q_4$ is constrained to be the intersection of the plane vertical to the above line of intersection and the ray-plane formed by the third side of $P_1P_2P_3P_4$. QED.

Consequently, there is only one 'virtual' rectangle (up to parallel displacement) that can share three sides of the retinal projection with a planar nonrectangular parallelogram. This rectangle will vary systematically as a function of the internal angle of the parallelogram and as a function of viewing angle.

Appendix 3. Derivation of the tilt of the virtual edge

The unique virtual rectangle that is consistent with the perceptual assumptions of model 1 (as shown in Appendix 2.2) will have a virtual leading edge AC, which is at 90° to the backboard (in the $x-z$ plane), and has the same projection as line AB in the image plane (ie $A'B' = A'C'$). In this appendix, we derive v, the tilt in 'world coordinates' of the virtual edge AC. Though the calculations in this appendix could be done in 3-D space, we have done them in 2-D space (horizontal and depth), using the simplification that for a vertical backplane (or building), all planes vertical to the long edges are also vertical.

The virtual leading edge AC lies in a plane that is vertical in world coordinates. The point at which the line of sight passing through B intersects this plane will give the coordinates of C (x_C, y_C, z_C). Since the plane is vertical, the x_C and z_C coordinates of the intersection point can be calculated in the $x-z$ plane (figure A3.1). In this figure, the point O represents the origin (focal point), AB is the leading edge of the parallelogram stimulus, and AC is the leading edge of the virtual rectangle. AB is assumed to have unit length. Points O and A have the same coordinates for all conditions, but the coordinates of point B depend on the viewing angle ω and the physical tilt ψ.

In the $x-z$ plane:

$$O = (0, 0), \quad A = (0, z_A), \quad B = (\sin\omega\cos\psi, z_A - \cos\omega) . \tag{A3.1}$$

In experiment 1, we chose viewing conditions such that $\omega = 180° - \theta$, otherwise, ω was fixed at a value determined by the experiment (90°, 45°, or 135°).

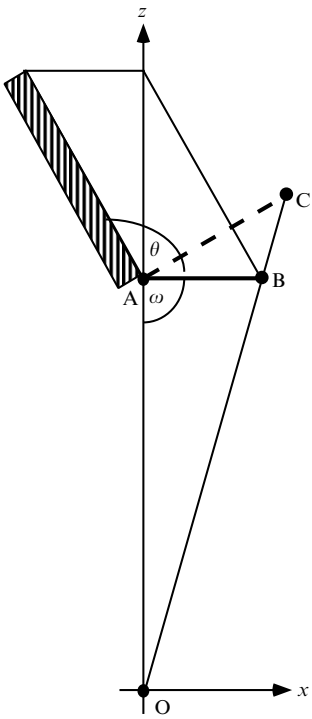

Figure A3.1. Plan (overhead) view of an example stimulus shape showing internal angle, θ, the leading edge of the stimulus, AB, the origin, O, from which the observer viewed the stimuli, and the angle between OA and AB, ω. The dashed line AC represents the leading edge of the inferred 'virtual' rectangle, and is at 90° to the shaded backboard.

In the $x-z$ plane, line AC is described by the equation

$$z = \tan(\omega + \theta - 180^\circ)\, x + z_A\,, \tag{A3.2}$$

and line OB by the equation

$$z = \frac{z_A - \cos\omega}{\cos\psi \sin\omega} x\,. \tag{A3.3}$$

Since point C is on the line joining O and B, it can be expressed as a linear combination of O and B, such that

$$(x_C,\, z_C) = t(x_B,\, z_B) + (1 - t)(0,\, 0)\,. \tag{A3.4}$$

Therefore

$$t = \frac{x_C}{x_B}. \tag{A3.5}$$

Lines AC and OB intersect at C. From equations (A3.2) and (A3.3)

$$\frac{z_A - \cos\omega}{\cos\psi \sin\omega} x_C = \tan(\omega + \theta - 180^\circ) x_C + z_A\,; \tag{A3.6}$$

therefore

$$x_C = \frac{z_A \cos\psi \sin\omega}{(z_A - \cos\omega) - \tan(\omega + \theta - 180^\circ)\cos\psi \sin\omega}\,. \tag{A3.7}$$

Substituting (A3.1) and (A3.7) into (A3.5), we obtain

$$t = \frac{z_A}{(z_A - \cos\omega) - \tan(180^\circ - \omega - \theta)\cos\psi \sin\omega}. \tag{A3.8}$$

In 3-D world coordinates, the point of intersection, C, is given by:

$$(x_C, y_C, z_C) = t(x_B, y_B, z_B) + (1 - t)(0, 0, 0) = t(\cos\psi \sin\omega, y_A + \sin\psi,\ z_A - \cos\omega) \tag{A3.9}$$

The vector between the point A(0, y_A, z_A) and point C is given by the difference vector

$$\overrightarrow{AC} = [t\cos\psi \sin\omega,\ t(y_A + \sin\psi) - y_A,\ t(z_A - \cos\omega) - z_A]\,. \tag{A3.10}$$

The angle between the y-axis (0, 1, 0), and AC can be derived from the scalar product of the two vectors:

$$v = \arccos\left[\left[\frac{t(y_A + \sin\psi) - y_A}{\{(t\cos\psi \sin\omega)^2 + [t(y_A + \sin\psi) - y_A]^2 + [t(z_A - \cos\omega) - z_A]^2\}^{1/2}}\right]\right], \tag{A3.11}$$

where t is computed from equation (A3.8). The tilt of vector AC relative to the horizontal is given by $90^\circ - v$.